Santa Fe Institute Working Paper 14-XX-XXX
arxiv.org:14XX.XXXX [cond-mat.stat-mech]

arXiv:1412.8520v1 [cond-mat.stat-mech] 30 Dec 2014

Understanding and Designing Complex Systems:

Response to "A framework for optimal high-level descriptions in science and engineering—preliminary report"

James P. Crutchfield,[1,*] Ryan G. James,[1,†] Sarah Marzen,[2,‡] and Dowman P. Varn[1,§]

[1] *Complexity Sciences Center and Department of Physics, University of California at Davis, One Shields Avenue, Davis, CA 95616*
[2] *Department of Physics and Redwood Center for Theoretical Neuroscience University of California at Berkeley, Berkeley, CA 94720-5800*

(Dated: December 31, 2014)

We recount recent history behind building compact models of nonlinear, complex processes and identifying their relevant macroscopic patterns or "macrostates". We give a synopsis of computational mechanics, predictive rate-distortion theory, and the role of information measures in monitoring model complexity and predictive performance. Computational mechanics provides a method to extract the optimal minimal predictive model for a given process. Rate-distortion theory provides methods for systematically approximating such models. We end by commenting on future prospects for developing a general framework that automatically discovers optimal compact models. As a response to the manuscript cited in the title above, this brief commentary corrects potentially misleading claims about its state space compression method and places it in a broader historical setting.

Keywords: information theory, rate-distortion theory, computational mechanics, information bottleneck, macrostates, microstates, statistical physics, coarse-graining, dimension reduction, minimum description length

PACS numbers: 02.50.-r 89.70.+c 05.45.Tp 02.50.Ey 02.50.Ga

I. INTRODUCTION

Building compact models of nonlinear processes goes to the heart of our understanding the complex world around us—a world replete with unanticipated, emergent patterns. Via discovery mechanisms that we do not yet understand well, we eventually do come to know many of these patterns, even if we have never seen them before. Such discoveries can be substantial. At a minimum, compact models that capture such emergent "macrostates" are essential tools in harnessing complex processes to useful ends. Most ambitiously, one would hope to automate the discovery process itself, providing an especially useful tool for the era of Big Data.

One key problem in the larger endeavor of pattern discovery is *dimension reduction*: reduce the high-dimensional state space of a stochastic dynamical system into smaller, more manageable models that nonetheless still capture the relevant dynamics. The study of complex systems always requires this. For better or worse, it is frequently accomplished in an *ad hoc* fashion [1, 2]. Indeed, it is desirable to have an overarching framework for this kind of analysis that can be applied across the many manifestations of complex systems, but to date such a broad theory has not been forthcoming. Thus, the need for this kind of research remains and is more timely than ever [3]. Not surprisingly, it has a long and active history.

This is the setting into which steps a recent `arxiv.org` preprint "A framework for optimal high-level descriptions in science and engineering—preliminary report" [4]. As a solution to the problem of dimension reduction, it advocates for *state space compression* (SSC): Form a compressed variable Y_t that predicts a target statistic Ω_t of a system's behavior $\ldots X_{t-1}X_tX_{t+1}\ldots$. When viewed in a historical context, it is unclear if SSC is more than an alternative notation and vocabulary for extant approaches to dimension reduction. Here, we explain this question. We are concerned about several instances in Ref. [4] where statements are made about research we either participated in or are quite familiar with that do not accurately reflect that work. The following comments air our concerns, providing several constructive suggestions.

Our response to Ref. [4] is organized as follows. We recall the history over the last half century driving interest and research on reconstructing optimal minimal

* chaos@ucdavis.edu
† rgjames@ucdavis.edu
‡ smarzen@berkeley.edu
§ dpv@complexmatter.org

models, specifically as it bears on nonlinear complex systems. We briefly recount the approach of computational mechanics, which defined what optimal predictive models are and gave the general solution to finding them. We then draw connections to predictive rate-distortion theory that systematically approximates those optimal models. We also comment on information-theoretic ways to quantify model complexity and predictive performance and how they trade-off against each other. Our goal is to respond directly and briefly to Ref. [4], but not to review the broad and extensive literature on the topic of optimal descriptions of complex systems. As such, citations are intentionally narrowed to support a single narrative thread.

II. RECONSTRUCTING LOW-DIMENSIONAL MODELS OF COMPLEX PROCESSES

The research program to automate theory building for complex systems has its origins in fluid turbulence—a high-dimensional system, if there ever was one—studied for many decades, of great practical import, and at the time, according to Heisenberg [5], one of the premier problems in nonlinear physics. Cracking this problem relied on developing a connection between the abstract theory of dynamical systems and hydrodynamic experiment. This came in the *attractor reconstruction* method that extracted "Geometry from a Time Series" [6] using one or a few signals from a high-dimensional complex system. Attractor reconstruction eventually led to demonstrating that deterministic chaos is the mechanism generating weak turbulence [7], verifying a long-standing conjecture [8] in dynamical systems theory that over-threw the decades-old quasi-periodic theory of turbulence due to Landau and Lifschitz [9].

The reconstruction method, though, only produced a reduced-dimension space of *effective states* of the infinite-dimensional fluid dynamics, ignoring the dynamical mechanism that generated the turbulent motion. Generalizing attractor reconstruction, Refs. [10, 11] introduced methods to infer the *effective theory* or equations of motion from time-series generated by chaotic dynamical systems. Reference [10], in particular, also provided a critique of the general reconstruction approach, highlighting its subjectivity—one must choose a class of model representation. Such *ad hoc* choices preclude an objective measure of a system's complexity. Reference [12] solved the subjectivity problem. It was the first to state the problem of optimal predictive, *minimal* models of complex processes and provided its solution—the ϵ-*machine* and its *causal states*. Using this foundation, it was able to define the *statistical complexity*—the first consistent, constructive measure of structural complexity. And, it introduced *intrinsic computation*—the idea that all physical systems store and process information. Influenced by the goals of artificial intelligence at the time, it also challenged future researchers to develop an *artificial science*—automating the construction of minimal causal models of complex systems. Using these methods, which came to be called *computational mechanics*, Ref. [13] went on to give the first quantitative definition of the emergence of macrostate organization in terms of increasing structural complexity and intrinsic computation. It argued that there is a natural hierarchy of causal organization and introduced a renormalization-group method for moving up the hierarchy—a method of genuine pattern discovery. A recent review of this history is found in Ref. [14].

III. CAUSAL STATES AND MACROSTATES

Identifying emergent organization, especially if pursued as a problem in theoretical physics, is often couched in terms of finding system *macrostates*. The metaphorical intuition behind this framing is roughly that emergent organization is the analog of the macroscopic, measurable properties of a thermodynamic system. A macrostate—say, given by particular values of temperature and pressure—is the set of the many microscopic molecular configurations or *microstates* that lead to the same measurable properties. In this framing, macrostates emerge from the microstates under the action of the dynamics on the microscopic scale as it relaxes to equilibrium [15]. The practicing statistical physicist, more specifically, often begins the analysis of a system's properties by searching for an "order parameter" or for insightful "coarse-grainings", which are analogous concepts at our general level of discussion here.

A system's causal states, being groupings of microscopic trajectories that capture a system's emergent behaviors and organization, play a role very analogous to the system's macrostates and not its microstates [16]. Reference [4] offers a view that is substantially at variance with this. SSC addresses this in terms of the descriptions to which given behavior is compressed. Notably, Ref. [4, page 6] starts by restricting the microscopic dynamics:

> The **microstate** ... evolves in time according to a (usually) Markovian process ...

And then, on page 41 it notes that the observations of the stochastic processes analyzed in computational me-

chanics:

> ... do not evolve according to a first-order Markov process, and so cannot be identified with the fine-grained values X_t of SSC. On the other hand, ..., evolution over the space of causal sets is first-order Markov. This suggests that we identify the causal states with SSC's fine-grained space, not its compressed space.

In the page 6 quote, however, SSC assumed first-order Markov dynamics on the underlying process (microstates). In the context described above, in which the target Ω_t is the future $X_t X_{t+1} \ldots$ and Y_t compresses the past $\ldots X_{t-2} X_{t-1}$, the causal states are interpreted as coarse-grainings *only* of X_t. Thus, in SSC the dynamics on the causal states and the observed process are *both first-order Markov*. These restrictions are unnecessary.

A telling consequence of SSC's misidentification of how computational mechanics is used is that the assumption of a first-order Markov process for the microstates simplifies the dimension-reduction problem to the point that the computational complexity of identifying causal states would fall in P and no longer be NP-hard, as it is more generally. Overall, the simplification to first-order Markov obscures the relationship between computational mechanics and SSC. Computational mechanics can be applied to any stochastic process, including the Markovian one assumed to govern SSC's microscopic variables $\ldots X_{t-1} X_t X_{t+1} \ldots$.

Apparently, despite being defined as a coarse-graining, SSC (incorrectly) associates causal states with its microstates X_t for no reason other than their Markovian nature. This is a confusing association for two reasons: First, the X_t values were introduced as being (usually) Markovian, not necessarily Markovian; and second, by construction computational mechanics' causal states are a coarse-graining of trajectories and so by definition are a form of state-space compression. SSC's association is based on reasoning that is fundamentally flawed, reflecting a basic misunderstanding of causal states and how to apply computational mechanics.

IV. STATE SPACE COMPRESSION VERSUS RATE DISTORTION THEORY

Now, let's turn to consider how SSC defines its "macrostates" via coarse-graining microstates. In this, we find a basic connection between SSC and Shannon's rate-distortion theory applied to prediction—what is called *predictive rate distortion* (PRD) [17]. SSC largely ignores this important connection to prior work. This is strange since Shannon introduced rate-distortion theory explicitly to solve dimension reduction problems [18, 19]: find a compact "encoding" of a data source subject to a set of constraints, such as transmission rate, accuracy, processing, delay, and the like. Optimal modeling can be framed in just this way, for example, as found in Rissanen's *minimum description length* approach [20]. The physics metaphor for building models extends to this setting, too: data are the microstates, compressed model variables are macrostates, and coding constraints are physical boundary conditions.

Take SSC's target Ω_t to be the future of X_t, and find some Y_t that compresses the past of X_t to retain expected accuracy in predicting Ω_t without unnecessarily increasing asymptotic coding cost of transmitting Y_t to another observer. When accuracy is heavily prioritized over coding cost, the causal states are recovered, as shown by Refs. [21, 22] and as discussed in Ref. [23]. In fact, the causal states are essentially the answer to the question: Among all possible compressions with minimal accuracy cost, which has the minimal coding cost? Both issues of accuracy and coding cost can be extended directly to the case in which the target Ω_t is some coarse-graining of the future of X_t, so that one searches for compressed predictive and perceptual features or macrostates. Recent PRD work shows, in fact, how to extract just such coarse-grained macrostates, given a process's ϵ-machine [17], noting that the latter can be calculated theoretically [24, 25] or estimated empirically [26–28].

However, a major conceptual difference between PRD and the SSC framework is that SSC explicitly restricts X_t to be first-order Markov; whereas, PRD can address arbitrary order and infinite-order Markov processes. The first-order Markov assumption is simply not realistic. Microstates are rarely experimentally accessible; e.g., the spike trains from neurophysiological studies are a very coarse-grained observable of underlying membrane voltage fluctuations [29] and, as such, their dynamics is often infinite-order Markov [30]. This remains true even if the underlying membrane voltage dynamics are first-order Markov.

PRD and computational mechanics actually do state space compression. They find compressed predictive representations of a time series. Moreover, computational mechanics finds a hidden Markov model (HMM) for the given (non-Markovian) time series. Indeed, that's the point of the ϵ-machine and, for lossy versions, the point of causal rate-distortion theory [17] and the recursive information bottleneck method [23].

V. CODING COST

This all leads to the central question of how to quantify the organization captured by these macrostates. PRD and computational mechanics use the statistical complexity—the Shannon information in the causal states or in their coarse-grainings. SSC takes issue with the use of information theory in these approaches [4, page 40]:

> Statistical complexity plays a role in the objective function of [77] that is *loosely analogous* to the role played by accuracy cost in the SSC objective function.

[Emphasis ours.] However, SSC's coding cost $H[Y_0]$ is the information contained in its version of the causal states. This is a simplified version of the statistical complexity C_μ, once one restricts to first-order Markov processes and predictive mappings from pasts $\ldots X_{t-2}X_{t-1}$ to Y_t. That is, C_μ is a *coding cost*, not an accuracy cost.

Moreover, on Ref. [4, page 40], we read:

> ... the authors consider a ratio of costs rather than a linear combination of them, like we consider here.

However, PRD objective functions—e.g., as described by Refs. [23, 31]—are linear combinations of mutual informations.

VI. MUTUAL INFORMATION AS ACCURACY COST

Reference [4, page 46] criticizes the use of general multivariate mutual information in noting a difficulty with systems governed by a time-varying dynamic:

> The underlying difficulty is inherent in the very notion of using mutual information to define accuracy cost, and is intrinsic to consideration of the relation between Y and X at more than two moments in time.

Using mutual information at more than two moments in time is simply not a problem for ergodic stationary processes [32]. Moreover, it's not necessarily a problem for nonergodic or nonstationary processes either. For these, several examples in Ref. [33], as well as those in Refs. [34, 35], calculate the past-future mutual information (*excess entropy*) that involves all moments in time.

However, there is perhaps another underlying difficulty with using mutual information as an accuracy cost: It is used alternatively as either the coding cost or an inverse accuracy cost in the information bottleneck method [36]. And, there, mutual information as an 'inverse accuracy cost' really amounts to an accuracy cost that employs the KL-divergence between $\Pr(\Omega_t|Y_t)$ and $\Pr(\Omega_t|X_t)$ in the more general PRD framework.

Generally, though, the rate-distortion theorem [37] justifies the use of mutual information as a coding cost regardless of distortion measure. And, helpfully, this has been extended to the nonasymptotically large block coding limit [38, 39], demonstrating that one should not be glib about introducing new coding costs.

VII. GENERALITY

SSC is offered as an improvement on the current literature for being a principled and constructive method of generating macrostates. Reference [4]'s abstract states:

> This State Space Compression framework makes it possible to solve for the optimal high-level description of a given dynamical system ....

This brings up two concerns, one explicit in the quote and one implicit.

First, what is provided is not a constructive framework; it does not provide methods to solve for the macrostates. Moreover, each of the provided macrostate examples is constructed in an *ad hoc* manner.

Second, the burden of proof lies with SSC's advocates. Since alternative frameworks exist and are contrasted against, at this date progress behooves the authors to provide examples where their framework succeeds and others fail.

SCC's lack in these regards should be compared with how alternatives had to develop new calculational methods for the challenging problems that are entailed in modeling complex systems. For example, computational mechanics extended methods from holomorphic functional calculus that now give closed-form expressions for a process's information measures [40].

We emphasize that SSC's $H(Y_0)$ is exactly the computational cost-based objective function used to identify causal states, in which we constrain ourselves to deterministic mappings $Y_t = f(X_t)$ such that $\Pr(\Omega_t|Y_{t_0}) = \Pr(\Omega_t|f(X_{t_0}))$. Indeed, more broadly interpreted, causal states are the unique and minimal macrostate choice for

SSC in the limit of high premium on accuracy and minimal concern about computational cost. Adapting the proof of Thm. 1 in Ref. [22] will be helpful here. In any case, the causal states are the minimal sufficient statistic for prediction [41]. In other words, any process statistic can be calculated from them and this raises the bar quite high for SSC's proposed alternative macrostates.

In terms of implementations, PRD and computational mechanics are constructive. In rate-distortion theory generally there's the Blahut-Arimoto algorithm and its generalizations for calculating coarse-grained macrostates [37, 42]. And, there are several statistical inference algorithms that estimate ϵ-machines from various kinds of data [26, 27, 43, 44].

The general problem of dimension reduction for complex systems is an important one, and we encourage efforts along these lines. We appreciate that the synopses of computational mechanics, PRD, and information measures above address but a small part SSC's stated goals and the goals earlier researchers have set. It is important, though, that SSC start with correct assumptions and an understanding of its antecedents. In any case, we hope that our comments remedy, in a constructive way, misleading impressions that Ref. [4] gives of the state of the art of computational mechanics and predictive rate distortion.

VIII. RUBBER, MEET ROAD

Finally, let's end on a practical note. We are advocates for a broad, even pluralistic approach to automated nonlinear model building—i.e., for artificial science. However, our repeated experience is that general "frameworks" seriously stub their toes on application to experiments. Despite this, we are still optimists. Those wishing to contribute, however, should pick at least one application area and drill all the way down to show their alternative discovers a particular new scientific phenomenon. This is a necessary calibrating exercise for any attempt at generality.

In contrast to SSC, its antecedents have done their due diligence. Rate-distortion theory, developed for over a decade into the information bottleneck method [36], has been applied to test how close sensory spike trains are to stimuli predictive information functions [45]. For our part, computational mechanics led to a new theory of structure in disordered materials [28] and to measuring novel information processing in neural spike trains [30]. Looking forward to future engineering of complex systems, the structural understanding developed in these applications moves us in a direction to design novel semiconducting materials for a new generation of computing technology and the next generation of nanoscale spike-train probes that will scale to monitoring the activity of thousands of neurons [46].

ACKNOWLEDGMENTS

The authors thank the Santa Fe Institute for its hospitality during visits and Chris Ellison, Dave Feldman, and John Mahoney for helpful comments. JPC is an SFI External Faculty member. This material is based upon work supported by, or in part by, the U.S. Army Research Laboratory and the U. S. Army Research Office under contracts W911NF-13-1-0390, W911NF-13-1-0340, and W911NF-12-1-0288. S.M. is a National Science Foundation Graduate Student Research Fellow and a U.C. Berkeley Chancellor's Fellow.

[1] A. N. Gorban, N. Kazantzis, I. G. Kevrekidis, H. C. Ottinger, and C. Theodoropoulos, editors. *Model Reduction and Coarse-Graining Approaches for Multiscale Phenomena*. Series in Complexity. Springer, 2006.

[2] W. H. A. Schilders, H. A. van der Vorst, and J. Rommes. *Model Order Reduction: Theory, Research Aspects and Applications*, volume 13 of *Mathematics in Industry*. Springer, 2008.

[3] J. P. Crutchfield. The dreams of theory. *WIRES Comp. Stat.*, 6(March/April):75–79, 2014.

[4] D. H. Wolpert, J. A. Grochow, E. Libby, and S. DeDeo. A framework for optimal high-level descriptions in science and engineering—Preliminary report, 2014. arxiv.org: 1409.7403.

[5] W. Heisenberg. Nonlinear problems in physics. *Physics Today*, 20:23–33, 1967.

[6] N. H. Packard, J. P. Crutchfield, J. D. Farmer, and R. S. Shaw. Geometry from a time series. *Phys. Rev. Let.*, 45:712, 1980.

[7] A. Brandstater, J. Swift, Harry L. Swinney, A. Wolf, J. D. Farmer, E. Jen, and J. P. Crutchfield. Low-dimensional chaos in a hydrodynamic system. *Phys. Rev. Lett.*, 51:1442, 1983.

[8] D. Ruelle and F. Takens. On the nature of turbulence. *Comm. Math. Phys.*, 20:167–192, 1971.

[9] L. D. Landau and E. M. Lifshitz, editors. *Fluid Mechanics*. Pergamon Press, Oxford, United Kingdom, 1959.

[10] J. P. Crutchfield and B. S. McNamara. Equations of motion from a data series. *Complex Systems*, 1:417 – 452, 1987.

[11] J. D. Farmer and J. Sidorowitch. Predicting chaotic time series. *Phys. Rev. Lett.*, 59:366, 1987.

[12] J. P. Crutchfield and K. Young. Inferring statistical complexity. *Phys. Rev. Let.*, 63:105–108, 1989.

[13] J. P. Crutchfield. The calculi of emergence: Computation, dynamics, and induction. *Physica D*, 75:11–54, 1994.

[14] J. P. Crutchfield. Between order and chaos. *Nature Physics*, 8(January):17–24, 2012.

[15] M. C. Cross and P. C. Hohenberg. Pattern formation outside of equilibrium. *Rev. Mod. Phys.*, 65(3):851–1112, 1993.

[16] We describe them as *mesoscopic states*—a structural level intermediate between raw system configurations (microstates) and the macrostates of thermodynamics.

[17] S. Marzen and J. P. Crutchfield. Circumventing the curse of dimensionality in prediction: Causal rate-distortion for infinite-order markov processes. 2014. SFI Working Paper 14-12-047; arxiv.org:1412.2859 [cond-mat.stat-mech].

[18] C. E. Shannon. A mathematical theory of communication. *Bell Sys. Tech. J.*, 27:379–423, 623–656, 1948.

[19] C. E. Shannon. Coding theorems for a discrete source with a fidelity criterion. *IRE National Convention Record, Part 4*, 7:142–163, 623–656, 1959.

[20] J. Rissanen. Stochastic complexity and modeling. *Ann. Statistics*, 14:1080, 1986.

[21] S. Still and J. P. Crutchfield. Structure or noise? 2007. Santa Fe Institute Working Paper 2007-08-020; arxiv.org physics.gen-ph/0708.0654.

[22] S. Still, J. P. Crutchfield, and C. J. Ellison. Optimal causal inference: Estimating stored information and approximating causal architecture. *CHAOS*, 20(3):037111, 2010.

[23] S. Still. Information bottleneck approach to predictive inference. *Entropy*, 16:968–989, 2014.

[24] J. P. Crutchfield and D. P. Feldman. Statistical complexity of simple one-dimensional spin systems. *Phys. Rev. E*, 55(2):R1239–R1243, 1997.

[25] D. P. Feldman and J. P. Crutchfield. Discovering noncritical organization: Statistical mechanical, information theoretic, and computational views of patterns in simple one-dimensional spin systems. 1998. Santa Fe Institute Working Paper 98-04-026.

[26] C. R. Shalizi, K. L. Shalizi, and J. P. Crutchfield. Pattern discovery in time series, Part I: Theory, algorithm, analysis, and convergence. 2002. Santa Fe Institute Working Paper 02-10-060; arXiv.org/abs/cs.LG/0210025.

[27] C. C. Strelioff and J. P. Crutchfield. Bayesian structural inference for hidden processes. *Phys. Rev. E*, 89:042119, 2014. Santa Fe Institute Working Paper 13-09-027, arXiv:1309.1392 [stat.ML].

[28] D. P. Varn and J. P. Crutchfield. Chaotic crystallography: How the physics of information reveals structural order in materials. *Curr. Opin. Chem. Eng.*, 7:47–56, 2015. Santa Fe Institute Working Paper 14-09-036; arxiv.org:1409.5930 [cond-mat].

[29] F. Rieke, D. Warland, R. de Ruyter van Steveninck, and W. Bialek. *Spikes: Exploring the Neural Code*. Bradford Book, New York, 1999.

[30] S. Marzen and J. P. Crutchfield. Time resolution dependence of continuous-time information measures. *in preparation*, 2014.

[31] C. R. Shalizi and J. P. Crutchfield. Information bottlenecks, causal states, and statistical relevance bases: How to represent relevant information in memoryless transduction. *Adv. Compl. Sys.*, 5(1):91–95, 2002.

[32] R. G. James, C. J. Ellison, and J. P. Crutchfield. Anatomy of a bit: Information in a time series observation. *CHAOS*, 21(3):037109, 2011.

[33] W. Bialek, I. Nemenman, and N. Tishby. Complexity through nonextensivity. *Physica A*, 302:89–99, 2001.

[34] N. Travers and J. P. Crutchfield. Infinite excess entropy processes with countable-state generators. *Entropy*, 16:1396–1413, 2014. SFI Working Paper 11-11-052; arxiv.org:1111.3393 [math.IT].

[35] M. Tchernookov and I. Nemenman. Predictive information in a nonequilibrium critical model. *J. Stat. Phys.*, 153:442–459, 2013.

[36] W. Bialek, I. Nemenman, and N. Tishby. Predictability, complexity, and learning. *Neural Computation*, 13:2409–2463, 2001.

[37] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Wiley-Interscience, New York, second edition, 2006.

[38] V. Kostina and S. Verdu. Fixed-length lossy compression in the finite blocklength regime. *IEEE Trans. Info. Th.*, 58(26):3309–3338, 2012.

[39] V. Kostina and S. Verdu. Lossy joint source-channel coding in the finite blocklength regime. *IEEE Trans. Info. Th.*, 59(5):2545–2574, 2013.

[40] J. P. Crutchfield, P. Riechers, and C. J. Ellison. Exact complexity: Spectral decomposition of intrinsic computation. submitted. Santa Fe Institute Working Paper 13-09-028; arXiv:1309.3792 [cond-mat.stat-mech].

[41] C. R. Shalizi and J. P. Crutchfield. Computational mechanics: Pattern and prediction, structure and simplicity. *J. Stat. Phys.*, 104:817–879, 2001.

[42] K. Rose. A mapping approach to rate-distortion computation and analysis. *IEEE Trans. Info. Th.*, 40(6):1939–1952, 1994.

[43] D. Pfau, N. Bartlett, and F. Wood. Probabilistic deterministic infinite automata. In *Adv. Neural Info. Proc. Sys.*, pages 1930–1938. MIT Press, 2011.

[44] D. P. Varn, G. S. Canright, and J. P. Crutchfield. ϵ-Machine spectral reconstruction theory: A direct method for inferring planar disorder and structure from X-ray diffraction studies. *Acta. Cryst. Sec. A*, 69(2):197–206, 2013.

[45] S. E. Palmer, O. Marre, M. J. Berry II, and W. Bialek. Predictive information in a sensory population. 2013. arXiv:1307.0225.

[46] A. P. Alivisatos, M. Chun, G. M. Church, R. J. Greenspan, M. L. Roukes, and R. Yuste. The brain activity map project and the challenge of functional connectomics. *Neuron*, 74:970–974, 2012.